

CLAIMS

What is claimed is:

1. A method for managing traffic flow through a multipath network, comprising:
 - forwarding a packet along a first link of the multipath network;
 - tracking a load of the first link subsequent to forwarding the packet; and
 - preserving the first link for a subsequent packet having the same flow address as the forwarded packet upon determining a desired load change of the first link is less than a predetermined value.
2. The method of claim 1, further comprising modifying link designations to forward packets along upon determining the desired load change of the first link is greater than the predetermined value.
3. The method of claim 2, wherein modifying link designations comprises designating a second link to send the subsequent packet along.
4. The method of claim 2, wherein modifying link designations comprises preserving the first link to send the subsequent packet along.
5. The method of claim 1, wherein tracking the load comprises tracking one or more variables associated with the load of the first link.
6. The method of claim 5, wherein the one or more variables comprise bandwidth of the first link.
7. The method of claim 5, wherein the one or more variables comprise a delay of the first link.
8. The method of claim 5, wherein the one or more variables comprise a loss rate of the first link.

9. A storage medium comprising program instructions that are executable by a processor and that cause the processor to:

adjust positions of one or more pointers used to partition traffic flow through a multipath network, wherein the positions of the one or more pointers are variable relative to a range of hash units that correspond to flow addresses within the multipath network.

10. The storage medium of claim 9, wherein the program instructions for adjusting the positions of the one or more pointers comprise program instructions for modifying a position of one pointer at a time.

11. The storage medium of claim 10, wherein the program instructions for adjusting the positions of the one or more pointers comprise program instructions for:

modifying a hash number of a first pointer positioned between a highest loaded link and a least loaded link; and
subsequently modifying a hash number of a second pointer positioned between a second highest loaded link and a second least loaded link.

12. The storage medium of claim 9, wherein the program instructions for adjusting the positions of the one or more pointers are directed for use by an individual router of the multipath network.

13. The storage medium of claim 12, wherein the program instructions for adjusting the positions of the one or more pointers comprise program instructions for:

calculating an average amount of load per hash unit for individual links coupled to the router; and
calculating a desired load change on the individual links.

14. The storage medium of claim 9, further comprising program instructions for selecting a link of the multipath network to send a packet along based upon a hash number representative of a flow address of the packet and relative hash numbers of one or more the pointers.

15. The storage medium of claim 14, further comprising program instructions for hashing the flow address of the packet.

16. A router, comprising:
multiple ports for coupling to links of a network; and
a storage medium comprising program instructions executable using a processor for
selectively directing a data packet to one of the multiple ports; and
altering one or more of the conditions by which the data packet is
selectively directed.

17. The router of claim 16, wherein the one or more conditions comprise hash number values of one or more variable pointers configured to partition a range of hash numbers associated with possible flow addresses of the data packet.

18. The router of claim 16, wherein the one or more conditions comprise specific hash number ranges associated with possible flow addresses of the data packet.

19. The router of claim 16, wherein the program instructions are executable using the processor for altering the one or more conditions to reflect a load balancing policy of the router

20. The router of claim 16, wherein the program instructions are executable using the processor for accounting for the capacity of the links coupled to the multiple ports when the one or more conditions are altered.

21. The router of claim 16, wherein the program instructions are executable using the processor for altering the one or more conditions to monotonically balance loads between two of the multiple ports.

22. The router of claim 16, wherein the program instructions are executable using the processor for redirecting the data packet to another of the multiple ports upon detecting the one multiple port cannot accept the data packet.

23. A network, comprising:
multiple stations configured to send and receive data packets; and
a plurality of routers interposed between the multiple stations and interconnected by a mesh of links, wherein each router is configured to
selectively direct a first packet along a link coupled thereto in accordance with one or more variable pointers included within the router; and
record the status of the one or more variable pointers to direct a second packet having the same source and flow addresses as the first packet along the same link.

24. The network of claim 23, wherein each router is further configured to alter the positions of the one or more variable pointers.

25. The network of claim 23, wherein each router is configured to track the load conditions of the links coupled thereto.

26. The network of claim 23, wherein at least one router of the plurality of routers comprises a means for changing transmission control protocol connections among links of differing loss rates associated with the router.

27. A network device, comprising:
a router; and

a means for periodically changing transmission control connections among links of different loss rates which are coupled to the router.

28. The network device of claim 27, wherein the means for periodically changing the transmission control connections comprises:

a means for numbering the links in descending order of their long-term average loss rates; and

a means for modifying designations of hash numbers associated with the links such that $1/n$ of the flows switch from the lowest-loss link to the highest-loss link and the remaining flows switch from a higher-loss link to a lower-loss link.

29. The network of claim 28, wherein the means for modifying the designations of hash numbers is configured to modify the hash number designations at an interval that is slightly longer than the life spans of the majority of flows that traverse the network.